by transporting the rules or indirectly by using what we learned about building expert systems in general. The knowledge base includes information about the side effects of each of the drugs and about the physiological mechanisms of these side effects. This information allows us to predict drug interactions and the likelihood of occurrence of various side effects in a given patient, and to base explanations on knowledge of the underlying physiology. The knowledge base also includes specific information about drug regimens, about preventing and treating side effects, and about how to take all of this into account in selecting a drug and dosage regimen for the individual patient.

D. List of Relevant Publications

1. Feinberg, M. and Lindsay, R. K.: *Expert systems in Psychiatry and Psychopharmacology*. Psychopharmacol. Bull., 22, 1986, 311-316.

2. Lindsay, R. K.: Expert Systems in Psychiatric Diagnosis: Rule-Based Systems. Presented at MedInfo86, Washington, D.C.

3. Feinberg, M.: What Psychiatrists Can't Do. Presented at Medinfo86, Washington, D. C.

E. Funding Support

None.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaboration and Program Dissemination via SUMEX

We have established via SUMEX a community of researchers who are interested in AI applications in psychiatry. We also have used the message system to communicate with other AI scientists at SUMEX and elsewhere.

B. Sharing and Collaboration with other SUMEX-AIM Projects

During this past year we have had no occasion to engage in collaboration with other SUMEX-AIM Projects.

C. Critique of Resource Management

Our sole use of the system this year has been for communication. This has been very useful, but hampered by difficulties in matching the characteristics of various networks and terminals. This has made use of SUMEX, even for mail, awkward. It would be helpful to have some assistance with these problems.

III. RESEARCH PLAN

A. Project Goals and Plans

Our immediate objective is to develop expert systems that can differentiate patients with the various subtypes of depressive disorder, and prescribe appropriate treatment. This system should perform at about the level of a board-certified psychiatrist, i.e. better than an average resident but not as well as a human expert in depression. Eventually, we plan to enlarge the knowledge base so that the expert system can diagnose and prescribe for a wider range of psychiatric patients, particularly those with illnesses that are likely to respond to psychopharmacological agents. We will design the system so that it could be used by non-medical clinicians or by non-psychiatrist M.D.'s as an adjunct to consultation with a human expert. We plan also to focus on problems of the user interface and the integration of this system with other databases.

B. Justification and Requirements for Continued SUMEX use

The access to SUMEX resources is essentially our sole means of maintaining contact with the community of researchers working on applications of AI in medicine. Although we have moved our system to local workstations, the communications capability of SUMEX will continue to be important.

We anticipate that our requirements for computing time and file space will continue at about the same low level for the next year.

C. Needs and Plans for Other Computing Resources

We anticipate that the need for additional computing power will continue to be met by local workstations.

D. Recommendations for Future Community and Resource Development

Valuable as the present SUMEX facilities are to us, they are in many ways limited and awkward to use. The major limitation we feel is the difficulty and sometimes the impossibility of making contact with everyone who could be of value to us. We hope that greater emphasis will be put on internetwork gateways. It is important not only to establish more of these, but to develop consistent and convenient standards for electronic mail, electronic file transfers, graphic information transfer, national archives and data bases, and personal filing and retrieval (categorization) systems. The present state of the art feels quite limiting, now that the basic concepts of computer networking have become available and have proved their potential.

We expect that the role of the SUMEX-AIM resource will continue to evolve in the direction of increased importance of communication, including graphical information, electronic dissemination of preprints, and database and program access. The need for computer cycles on a large mainframe will diminish. We hope to have continued access to the system for communication, but do not anticipate continued use of it as a Lisp computation server.

If fees for using SUMEX resources were imposed, this would have a drastically limiting effect on the value of the system to us. Even if we had a budget to purchase such services, the inhibiting effect of having a meter running would cause us to make less use of it that we should. We have been conscious of the costs of the system and feel that we have not used it imprudently, even though we have not directly borne its costs.

IV.D.3. Dynamic Systems Project

Decision Support for Time-Varying Clinical Problems

Lawrence Widman, M.D., Ph.D.
Division of Cardiology
Case Western Reserve University
2065 Adelbert Road
Cleveland, OH 44106
(216) 844-3153

I. SUMMARY OF RESEARCH PROGRAM

A. PROJECT RATIONALE

Time-varying systems, which include many areas of medicine, science, economics, and business, can be described mathematically by differential equations. They are distinct from the pattern-matching and logic-based domains dealt with so successfully by existing expert system methods, because they can include feedback relationships. It is generally felt that they are best approached by enhancement of existing methods for deep model-based reasoning.

The goal of this project is to develop AI methods for capturing and using knowledge about time-varying systems. The strategy is to address general problems in model-based knowledge representation and reasoning. The intermediate objective is to develop methods which are powerful enough to work in selected realistic situations yet are general enough to be transportable to other, unrelated knowledge domains.

The tactical approach is to work on well-defined yet complex and interesting problems in the medical domain. We have, therefore, selected the human cardiovascular system as our prototype of a time-varying system, and are developing methods for representing and reasoning about its mechanical and electrical activities in the normal and diseased states.

A.1 Technical Goals

This project presently has two distinct tracks: hemodynamic modeling and cardiac arrhythmia interpretation.

1. Hemodynamic Modeling

The goals of this subproject are to develop:

(a) a knowledge-representation method using symbolic modeling which captures the qualitative and, when possible, the quantitative behavior of systems with feedback relationships. Preferably, the symbolic model should be translatable into the differential equations which describe the behavior of the system being modeled.

(b) a reasoning method based on the symbolic modeling tool created in subgoal (a) which permits the inference of differential diagnoses (a set of hypothesized diagnoses) from incomplete data.

(c) a reasoning method based on subgoals (a) and (b) which permits inference of the state of the model for each hypothesized diagnosis. This subgoal would be satisfied by an algorithm which specifies a self-consistent set of values for all variables in the model, for a given hypothesis based on a given set of data. Such sets of data would constitute initial conditions for differential equations derived from the model.

(d) a simulation method, based on the model and its equivalent differential equations together with the initial conditions derived from the differential diagnosis (steps a-c above), for predicting the expected time course of the system being modeled for each hypothesized diagnosis. This method could also be used to predict the effects of treatments being considered for recommendation by the program.

(e) a reasoning method, based on domain-independent properties of the model, for shrinking and/or expanding the model automatically to use a minimal model configuration to account for normal and abnormal data.

(f) an explanation facility for examining the model, the given data, the inferred hypothesized diagnoses, predicted behaviors, and modifications of the model, to answer user queries and to teach fundamental concepts.

2. Cardiac Arrhythmia Recognition

The goals of this subproject are to develop:

(a) a symbolic model of the electrical system of the human heart, including pertinent anatomic and electrophysiologic features of the normal and diseased heart. The electrophysiologic features would include deterministic characteristics (e.g., conduction velocities, refractory periods), stochastic features (e.g., behavior of automatic foci), and temporal interactions (e.g., competing pacemakers).

(b) a symbolic/numeric representation of the observable features of the electrical activity of the heart, both surface EKG and intracardiac recordings, including noise. This representation would be intended to allow a feature extraction module working on actual patient data to communicate with a symbolic reasoning module, and would be translatable directly into waveform display format.

(c) a reasoning method for extracting features from raw, digitized signal data. This method would augment established signal processing techniques by using knowledge-based algorithms to improve detection of P and T-U waves and to improve rejection of noise. It should be noted that this is itself a major research undertaking in the signal processing domain.

(d) a reasoning method for inferring the cardiac rhythms consistent with a given disease state in the model, similar to the prediction of consequences of the hemodynamic model in the first subproject. The output of this method would be in the symbolic/numerical representation of subgoal (b).

(e) a reasoning method for inferring possible disease states in the model from a given feature-extracted recording of the electrical activity of the heart. This subgoal constitutes cardiac arrhythmia interpretation, and is itself a major research project.

(f) a categorization method for inferring hierarchies of diagnoses from elementary abnormalities. For example, "periods of atrial fibrillation up to 30 minutes at up to 150 beats/min, supraventricular tachycardia of up to 10 beats length at a rate of 130 beats/min, and sinus bradycardia with a minimum rate of 45, all consistent with the sick sinus ("tachy-brady") syndrome" and "two QRS morphologies are present: they are narrow at rates less than 120 and are wide at rates above 120, consistent with a rate-dependent bundle branch block".

(g) an explanation facility for examining the model, the input data, and the interpretations to answer user queries and to teach fundamental concepts.

B. MEDICAL RELEVANCE AND COLLABORATIONS

The two subprojects have related but separate medical goals:

1. Hemodynamic Modeling.

There are three subgoals in this subproject: model-based sensor integration, model-based caregiver assistance, and model-based experiment interpretation.

a. Model-based Sensor Integration.

The long-range application of this subproject is the integration of patient-related data in the intensive care environment. Model- based real-time systems would allow the system to share a global understanding of the patient's condition with the human caregivers. Thus, it could interpret significant trends in key parameters and could draw attention to relationships which might otherwise escape attention in the constant flood of data common to these environments.

b. Model-based Caregiver Assistance.

It could also serve as an assistant to the caregiver. In this mode, the human caregiver could evaluate the merits of proposed diagnostic and therapeutic measures in light of available data on the patient's condition.

Practical application of these concepts requires further development of the model and the reasoning algorithms, and extensive testing against real clinical scenarios. Refinement and quality control are presently the responsibility of the principal investigator, who is a board-certified internist with subspecialty training in invasive cardiology.

Practical application also awaits gen . acceptance of standardized hospital data buses for automatic acquisition of importa... parameters now stored primarily on paper or on computers outside the intensive care setting, such as fluid inputs and outputs, medications, and results of invasive and non-invasive tests. Further, improved user interfaces will require better graphics and increased computer literacy on the part of caregivers.

c. Model-based Experiment Interpretation.

An intriguing third application of this subproject is in the area of interpretation of biomedical experiments. The symbolic model concept, which enforces objectivity, can assist investigators by allowing them to compare alternate interpretations of experimental data. In this application, several alternate models would be proposed by the experimenter to explain a given experimental outcome. The consequences of each model given different experimental parameters could then be evaluated and compared with real data to confirm or refute competing proposed models.

The advantage of using a computer in this manner is the guarantee of self-consistent and objective exploration of each possibility. The advantage of using a symbolic model, rather than a numerical model such as the Guyton-Coleman model or a simpler derivative, is that the underlying cause and effect relationships are explicit and can be easily modified by the experimenter. The AI interest in this subgoal would be the refinement of the symbolic model through application to real experiments. Unlike in the MOLGEN project at Stanford, automatic hypothesis formation would not be an objective in this subgoal.

A new collaboration to explore this application is being explored with Dr. E. Merrill Adams, an experimental physiologist in the Department of Surgery at Case Western Reserve University School of Medicine. Dr. Adams approached us because of his long-standing interest in applying AI techniques to his experiments on the interactions of the cardiovascular and pulmonary systems. A discussion group is being organized, and we hope to continue despite the move of the principal investigator to Texas this summer.

2. Cardiac Arrhythmia Recognition.

The long-range application of this sub-project is in clinical devices such as intensive-care arrhythmia monitors, portable Holter monitors, and implantable cardioverter-defibrillators. There are two subgoals: recognition of surface electrocardiographic (EKG) recordings and recognition of intracardiac recordings.

a. Recognition of surface electrocardiographic recordings.

Substantial and well-recognized obstacles in signal processing will likely prevent non-AI algorithms from advancing beyond the current state of the art of interpretation of surface EKG recordings. These obstacles are primarily the problems of reliable detection of P and T-U waves, and rejection of noise. We hope that AI techniques will be helpful with these problems, as is suggested by the work of Muldrow et al. (Computers and Cardiology, 1986, in press).

We hope further that, by mimicking the behavior of expert human cardiologists, these obstacles can be bypassed if they cannot be overcome. We have enlisted as consultants Dr. William Long of M.I.T., who supervised Muldrow in the paper cited above, and Dr. Benjamin Kuipers of the University of Texas at Austin, who is interested in AI techniques for physiological modeling.

b. Recognition of intracardiac recordings.

Intracardiac recordings, which are taken from wires placed in the heart by percutaneous venous puncture or around the heart by surgery, are relatively free of P wave ambiguity and of noise. They are representative of the quality of signals available to implantable cardioverter-defibrillators.

Cardioverter-defibrillators are devices like pacemakers in that they monitor the heart rhythm in a patient to determine if an abnormality exists. They are capable of taking action (electrical countershock) if an appropriate abnormality is detected. Unlike ordinary pacemakers, these devices detect abnormalities characterized by rapid rates of heart activity, rather than excessively slow rates. They have been shown to reduce one-year mortality in high-risk patients from 30% to 2%, and they are expected to play an increasingly large role in treatment of such patients.

These relatively new devices currently use quite simple algorithms to detect abnormalities. The action they take consists of applying an electrical shock directly to the heart. This shock is frequently unpleasant to the patient. The problem is that the algorithms sometimes confuse innocent rapid heart rates, such as from exercise or atrial tachyarrhythmias, with lethal ventricular arrhythmias. This has proved troublesome enough to prompt repeated calls in the electrophysiology literature for improved algorithms for arrhythmia recognition in these devices.

The algorithms developed in this subproject would be suitable for this application when the computer power in the devices improves. Because these devices require powerful energy sources to perform repeated shocks over their lifetimes of 2-3 years, the power drain of more sophisticated computer chips is less important than it would be in ordinary pacemakers.

C. Highlights of Research Progress

1. Hemodynamic Modeling

Subgoals (a) through (d) have been accomplished in prototype form. The approach relies on a semi-quantitative representation [subgoal (a)] which assigns values by default if the user does not specify more detailed information. The second phase of this project yielded subgoal (d), the simulation of a given model. This phase was accomplished by translating the model into a set of dynamical systems equations, which were then integrated in the standard manner.

More recently, subgoals (b) and (c) have been accomplished in prototype form. Constraint propagation using a dynamically generated semi-quantitative quantity space is performed by interpreting the model as a set of constraint equations. Domain-independent heuristics which recognize morphological features of the model are used to further constrain the propagation of constraints and to generate hypotheses when ambiguities arise. These heuristics generate a set of self-consistent hypotheses, each of which is a hypothesized diagnosis (subgoal b). Dr. Yong-Bok Lee of the Case Western Reserve University Department of Electrical Engineering and Applied Physics participated in this subgoal for his doctoral dissertation. The doctoral dissertation, awarded in August, 1986, was co-supervised by Professor Yoh-Han Pao of that Department and by Dr. Widman.

Each hypothesized diagnosis is then refined by mathematical relaxation, in which the propagated values are treated as initial guesses, and the values are refined iteratively, again by interpreting the model as a set of constraint equations (subgoal c). In the several scenarios which have been examined, the value assignments achieved by hypothesis and iterative refinement have achieved correlation coefficients up to 0.90 with the values obtained by simulation of the same model.

We do not anticipate beginning work on the remaining subgoals until the above prototype methods have been further refined and tested.

2. Cardiac Arrhythmia Recognition

This subproject is just beginning. We have built a prototype symbolic model of the electrical conduction system of the heart and have reproduced simple rhythms. The important issues of stochastic variation and of noise have not been addressed. We are hopeful that important insights will obtained from newly developing literature on stochastic simulation (e.g., Pearl, J.: *Evidential Reasoning Using Stochastic Simulation of Causal Models*. Artificial Intelligence. 1987;32:245-257).

Following the move of the principal investigator to Texas, this subproject will replace the hemodynamic modeling subproject as the major research focus. This research effort will be supported in part by a Grant-in-Aid from the American Heart Association, Texas Affiliate.

The principal investigator will have access to intracardiac signals from a variety of appropriate patients on the clinical service at his hospital complex. This should facilitate the development of practical algorithms.

We have also begun discussions with a major pacemaker manufacturer with the goal of establishing a working relationship. The purpose of the relationship would be to enable practical pacemaker manufacturing constraints to be taken into account from an early stage in the development of this subproject. So far, the discussions have demonstrated interest on both sides, but will require further algorithm development in order to proceed.

D. List of Relevant Publications

1. Widman, L.E. Reasoning about Diagnosis and Treatment in a Causal Medical Model using Semi-Quantitative Simulation and Inference. Workshop on Artificial Intelligence in Medicine, National Conference on Artificial Intelligence AAAI-87, Seattle.

2. Widman, L.E., Lee, Y.-B., and Y.-H. Pao. Diagnosis of Causal Models by Semi-Quantitative Reasoning. (submitted to SCAMC 1987).

3. Widman, L.E., Lee, Y.-B., and Y.-H. Pao. Diagnosis of Causal Medical Models by Semi-Quantitative Reasoning. In: Miller, P.L. (ed.). Topics in Medical Artificial Intelligence, Springer-Verlag (in preparation).

4. Lee, Y.-B. and L.E. Widman. Reasoning about Diagnosis and Treatment in a Causal Time-varying Domain using Semi-Quantitative Simulation and Inference. Workshop on Artificial Intelligence and Simulation, National Conference on Artificial Intelligence AAAI-86, Philadelphia.

5. Widman, L.E. Representation Method for Dynamic Causal Knowledge Using Semi-Quantitative Simulation. Fifth World Conference on Medical Informatics. 1986: 180-184.

E. Funding Support

1. American Heart Association, Texas Affiliate
Grant-in-Aid Award.
Knowledge-Based Computer Algorithms for Arrhythmia Analysis.
Principal Investigator: Lawrence E. Widman.
Award period: July, 1987 - June, 1988.
Level: $24,850 direct costs.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Sharing and interactions with other SUMEX-AIM projects

The major interactions with SUMEX-AIM have been (1) computational support and (2) communication with members of the AIM community.

(1) SUMEX-AIM is the major source of computing power at this time. Dr. Widman expects that a LISP workstation will be available after he relocates to Texas this summer. SUMEX-AIM computing power will then be needed primarily for demonstrations at meetings and as backup during workstation down-time.

(2) SUMEX-AIM is the current electronic mailbox. Its central location allows ready Email access by users of Arpanet, Bitnet and Csnet. This access has proved invaluable to Dr. Widman in communicating rapidly and effectively with co-workers at other institutions. The value of this type of communication has been demonstrated several times during the past year, when he had to make major career, equipment negotiation, and manuscript revision decisions, without local expertise, within short periods of time.

Review of the longer term history of this project shows that it would not exist had SUMEX-AIM not provided telecommunication support for the initial feasibility project in 1984-1985, which was carried out on the computers of the MIT Laboratory for Computer Science, Clinical Decision Making Group.

C. Critique of Resource Management

The service provided by SUMEX-AIM has been exemplary, largely because of prompt and effective response to difficulties as they arise. There has been a clear effort to assure that telecommunication access remained reliable during changes in commercial vendors, and the staff have responded to several technical questions promptly and accurately. Down-time has been minimal compared to that of other systems we have used, and is almost always scheduled several days in advance.

The reason we sought contact with the AIM community was that it seemed the natural niche for our research interests. There is no short-term prospect that this project will reach commercial maturity or that it will lose sight of fundamental AI issues, and so we feel that it still belongs in the scientific AIM framework.

As noted in the previous section, the communication with other members of the AIM community has proved invaluable in the advancement of this project.

III. RESEARCH PLANS

Project Goals and Plans

The long range goals of this project are to develop intelligent comprehensive monitoring/alarm systems for intensive care unit settings; and intelligent arrhythmia recognition systems for monitors, Holter recorders, and implantable cardioverter/ defibrillators. The short term strategies for achieving these goals are discussed above.

The next phase of this research will be conducted at the University of Texas Health Science Center at San Antonio. Dr. Widman will be joining the faculty of Medicine there on July 1, 1987, in the Division of Cardiology. His clinical duties will include invasive hemodynamic and electrophysiological studies on selected patients. Substantial time is committed to research, and this project will constitute his major research emphasis.

B. Justification and requirements for continued SUMEX use

The justification for this project is its potential for advancing the state of the art of expert system technology in the area of temporal reasoning and deep causal modeling, and for demonstrating practical use of expert symbolic computing in potentially life-saving, knowledge-intensive environments.

The requirements for continued SUMEX-AIM use should be the same as currently: telecommunications support, Arpanet access, about 3 megabytes of disc space, and a reasonable amount of CPU time. When the Lisp workstation becomes available (see below), the requirement for telecommunication support and CPU time should decrease.

C. Needs and plans for Other Computing Resources Beyond SUMEX-AIM

The symbolic computing needs for the hemodynamic modeling subproject are being met by SUMEX. Once embarked on the arrhythmia recognition subproject, there will be a strong need for high-resolution graphics, and processing of tens of megabytes of data. To meet these needs, a Lisp workstation will be provided by the University of Texas. Data acquisition in real time and initial signal processing will be done with an IBM AT class microcomputer equipped with a standard third-party multichannel analog-to-digital converter. Communication between the machines will be by RS232 or the local Ethernet LAN. Once these machines are in place, SUMEX-AIM will be needed primarily for communication and demonstration projects, as noted above.

D. Recommendations for Future Community and Resource Development

Our strong recommendation is that SUMEX-AIM be maintained as a national AIM resource for communication, development of software useful to the AIM community, and sharing of demonstration projects. SUMEX-AIM could also serve as a central source of advice for new workstation users who may be geographically isolated from experienced workstation users.

Additionally, we would strongly support retention of the current telecommunication support and enough computing power to support promising young investigators who would otherwise not have access to symbolic computing power.

IV.D.4. Knowledge Engineering for Radiation Therapy

KNOWLEDGE ENGINEERING FOR RADIATION THERAPY

Ira J. Kalet, Ph.D.
Witold Paluszinski
University of Washington
Seattle, Washington

I. Summary of Research Program

A. Project Rationale

We are developing an expert system for planning of radiation therapy for head and neck cancers. The project will ultimately combine knowledge-based planning with numerical simulation of the radiation treatments. The numerical simulation is needed in order to determine if the proposed treatment will conform to the goals of the plan (required tumor dose, limiting dose to critical organs). The space of possible radiation treatments is numerically very large, making traditional search techniques impractical. Yet, with modern radiation therapy equipment, the design of treatment plans might be significantly aided by automatically generating plans that meet the treatment constraints. The project will result in systematization of knowledge about radiation treatment design, and will also provide an example of how to represent and solve design problems with a knowledge based system.

B. Medical Relevance and Collaborations

Radiation therapy has shown dramatic improvement in the cure rate for many tumor sites in the last two decades. Much of this can be attributed to the improved penetration capability of modern megavoltage X-ray machines. These high energy beams can deliver high tumor doses without overdosing surrounding tissue in many cases. However, they are typically used in very limited ways, because of the lack of suitable simulation systems to compute the dose distribution for any but a few narrow choices of treatment geometry. In the last few years these simulation systems have been extended to the full range of geometric treatment arrangement that any therapy machine is capable of. Thus it would be valuable to be able to generalize our knowledge of treatment technique by exploring these expanded possibilities. In addition, even treatments with standard geometries can be very complex, and it is tedious to explore all of them individually. A knowledge-based system can generate a few "best" plans which satisfy the constraints and allow more time for the physician to evaluate the options, or make minor adjustments for optimization.

Since cancer treatment is a multi-disciplinary approach involving surgery and chemotherapy as well as radiation, it is important to coordinate this work with knowledge-based program projects in those areas. Most significant is the ONCOCIN project, which addresses management of patients on chemotherapy protocols.

This project has some relevance to computer science as well, in that our approach, if successful, may contribute to a better understanding of design problem solving with knowledge-based systems.

C. Highlights of Research Progress

In the past year, we have made significant additions to the rule database for details of head and neck cancer treatment. We have devised a representation of parameters for radiation treatment fields and created a set of prototype treatment field arrangements. The prototypes are used as building blocks for constructing complex treatment plans. In addition we have examined the issues of control strategy associated with using prototypes in planning.

Our expert system now has about two hundred rules, a two-level (agenda-based) control strategy, and about ten prototypes for plan construction. It is written in Interlisp on a VAX running the VMS operating system. This environment was chosen because it is also the environment used for a graphic simulation system that does radiation dose calculations for arbitrary treatment plans. The dose calculation is needed to determine whether a plan meets the treatment goals set by the system in its early phases of planning.

D. List of Relevant Publications

1. I. Kalet and W. Paluszynski: A Production Expert System for Radiation Therapy Planning. Proceedings of the AAMSI Congress 1985, May 20-22, 1985, San Francisco, California. Edited by Allan H. Levy and Ben T. Williams. American Association for Medical Systems and Informatics, Washington, D.C., 1985.

2. W. Paluszynski and I. Kalet: Radiation Therapy Planning: A Design Oriented Expert System. WESTEX-87 (Western Conference on Expert Systems), Anaheim, California, June 2-4, 1987.

3. I. Kalet and J. Jacky: Knowledge-based Computer Simulation for Radiation Therapy Planning. Proceedings of the Ninth International Conference on the use of Computers in Radiotherapy, Scheveningen, the Netherlands, June 1987. North Holland, 1987.

II. Interactions with the SUMEX-AIM Resource

Our main use of the SUMEX-AIM resource has been as a means to be in contact with other researchers working on AIM projects. The existence of a mailbox at SUMEX-AIM has made it much easier for colleagues at other institutions to communicate with us, and has been valuable in assisting us with organizing the AIM Workshop for 1987.

We have had a great deal of contact with members of the ONCOCIN project and other groups. This has been valuable to us in stimulating creative approaches to our project.

III. Research Plan

A. Project Goals and Plans

We plan to continue to acquire rules and develop our current expert system. This includes solving problems of use of prototypes, satisfaction of constraints by some kind of backtracking search, and incorporating evaluation of plans by using the results of the dose computation. This last idea involves coupling the expert system with the dose computation system (written in PASCAL) in suitably efficient ways. Our long-term goal is to shape the user interface and improve the system performance to where it can provide assistance to clinicians in treatment design for patients in the normal course of treatment.

B. Justification and Requirements for Continued SUMEX use

We foresee continued need to be in touch with other members of the AIM community, particularly projects centered at SUMEX. While we do not expect to use the computing resources of SUMEX directly, some more extensive communication and involvement is likely to be useful.

C. Plans For Other Computing Resources

The main computing resources for our project will continue to be local. We will be rewriting the expert system code in VAX Lisp, an implementation of Common Lisp on the DEC VAXstation. We expect delivery of a VAXstation II/GPX in the near future. This appears to be a good choice to satisfy our need for high performance graphic simulation and a reasonable Lisp system. However, the resources for the dose computation may not be adequate as we incorporate more sophisticated computation models. As this develops, we hope to experiment with distributed systems, in which the dose computation may run on a remote resource, which may or may not be at SUMEX.

D. Recommendations for Future Community and Resource Development

Two areas will be of increasing importance to us in the future: communication capabilities (electronic mail and file transfer) and centralized databases. By centralized databases, we refer to the need for better maintenance of mailing lists, information about projects, and possibly on-line reports. Dr. Kalet's experience in organizing the AIM Workshop for 1987 demonstrated that electronic communication is invaluable, even in its present state, but in order to create a list to send announcements to, we expended many hours of manually cutting and pasting messages containing past lists and searching for up-to-date electronic mail addresses.

If fees for use of SUMEX resources were imposed, the main impact on our project would be one of increased isolation, unless we could find grant support for the fees.

IV.D.5. Pathophysiologic Diagnosis Project

COMPUTER-BASED EXERCISES IN PATHOPHYSIOLOGIC DIAGNOSIS

J. Robert Beck, M.D.
Dartmouth College School of Medicine
2 Maynard St.
Hanover, N.H. 07355

I. SUMMARY OF RESEARCH PROGRAM

A. Project rationale

Research in artificial intelligence at Dartmouth Medical School focuses on three main areas: 1) knowledge-based systems applied to laboratory medicine and pathology, 2) knowledge acquisition using machine learning techniques, and 3) computer-based instruction using artificial intelligence techniques to critique students' workup plans. These projects have in common the fundamental research questions of how knowledge should be represented and used in a classification approach to problem-solving related to the use of laboratory data.

Knowledge-based systems in laboratory medicine:

We are investigating the use of knowledge-based systems to review requests for blood products. A system is being developed to advise pathologists and pathology residents about the appropriateness of transfusion requests.

The system will have both diagnostic and therapeutic objectives. The diagnostic part of the system will be used to evaluate information available in machine-readable form, and then ask the user a few relevant questions. Based on the available information, the system will determine possible diagnoses relevant to transfusion medicine. The current prototype is focusing on coagulopathies and bleeding disorders. The objective is to have a system that can quickly provide a summary of relevant laboratory information to the pathologist charged with the responsibility of evaluating appropriateness of transfusion requests. The therapeutic recommendations of the system will be focused on determining appropriate choices and quantities of blood products or substitutes. One of the purposes of this investigation is to determine whether a knowledge-based system can eventually reduce inappropriate use of blood products. The purpose of the tool is not to usurp decision-making, but to pre-process large volumes of transfusion requests and large volumes of data on each request, in order to focus the pathologist's attention in a time-efficient manner on the most relevant information. The system is in the early knowledge acquisition stage. The initial prototype is being built using IBM's Expert System Environment tool.

Knowledge acquisition for knowledge-based systems:

The purpose of this project is to develop machine learning tools that can be used for knowledge acquisition from databases. The focus is on deriving classification rules in the form of criteria tables. The criteria table format has been used for many years in medicine, and is still in use particularly in the area of rheumatic diseases (for example, the ARA criteria for systemic lupus erythematosus). Other diseases for which diagnostic criteria tables have been developed include polycythemia vera, multiple myeloma and primary biliary cirrhosis. In addition, criteria tables have been found useful as a knowledge representation for expert systems.

We have developed a program, called the CRiteria Learning System (CRLS) which is capable of automatically generating criteria tables from a database of positive and negative examples. CRLS is implemented in Common LISP on a SUN-3 workstation. It utilizes not only the raw data but also some background knowledge supplied by the user about the concepts to be learned, the features of the problem, and the type of diagnostic performance the user wishes to optimize (i.e. sensitivity, specificity, efficiency, etc.). CRLS learns decision rules that are more comprehensible than the rules generated by other machine learning programs. Tests of the system have also shown that it is capable of handling large databases containing as many as 1500 cases with 50 variables each.

Teaching medical pathophysiology using computer-based tools:

The project "Computer-based Exercises in Pathophysiologic Diagnosis" is funded through the National Library of Medicine's Medical Informatics research initiative. It has four specific aims:

1. To develop two computer-assisted laboratory exercises for basic content areas (anemia and coronary artery disease) in second-year medical education, oriented toward the processes of diagnosis and evaluation, utilizing techniques of medical decision science, critiquing, and software engineering (the PLAN-ALYZER system);

2. To utilize the computerized teaching modules to test the hypothesis that students with access to process-oriented educational tools can integrate their didactic knowledge more effectively than with access only to non-process oriented traditional education, including lecture notes, texts, and non-intelligent audiovisual aids;

3. To develop practical application versions of the two PLAN-ALYZER models that can be used as diagnostic tools with more senior medical students, residents, and physicians in the clinic, providing them with a decision analysis tool and expert critiques of their evaluations of real patients;

4. To utilize the originally proposed and the advanced systems to explore the process of how the effective physician solves clinical problems, a process which has been found to be different from traditional problem solving.

PLAN-ALYZER prototyping is being accomplished on the Macintosh Plus and Macintosh II workstations, using the Macintosh Programmer's Workshop. Novel AI features of the PLAN-ALYZERs include a scoring metric based on unate boolean functions, to compare students' decision trees with gold standard trees, a mechanism by which augmented transition network critiques can be developed for decision models, and the encoding of the domain experts' instructional styles as well as content into the models.

An interdisciplinary team of computer scientists, physicians, and educators is working on the Computer-based Exercises project. A prototype system is nearing completion, with formative evaluation scheduled for Fall, 1987.

D. Relevant Publications

Beck, J.R., Prietula, M.J., Russo, E.A.: A role for intelligent systems in teaching medical pathophysiology. In: Salamon, R., Blum, B., Jorgenson, M. (eds). Proc. Fifth Conf. Med. Inform. (MEDINFO '86), Elsevier-North Holland, Amsterdam, 1986, 936-938.

Beck, J.R.: Artificial intelligence: A topic for Medical Decision Making? (edit.) Med. Decis. Making 1987; 7:4.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

The Dartmouth group is pleased to be a new addition to the SUMEX research resource. Most of our projects take place on the Dartmouth campus, but we require access to the national AI community in order to share ideas, disseminate research results, and grant our trainees and junior faculty access to the developments of others. Also, inasmuch as our research in medical educational applications of computer science and decision making has significant potential for dissemination, the SUMEX community of scholars forms a natural group for focusing and broadening our research ideas.

Appendix A
AIM Management Committee Membership

Following are the current membership lists of the various SUMEX-AIM management committees:

AIM Executive Committee:

SHORTLIFFE, Edward H., M.D., Ph.D. (Chairman)
Principal Investigator - SUMEX
Medical School Office Building, Rm. X271
Stanford University Medical Center
Stanford, California 94305
(415) 723-6970

FEIGENBAUM, Edward A., Ph.D.
Co-Principal Investigator - SUMEX
Heuristic Programming Project
Department of Computer Science
701 Welch Road, Building C
Stanford University
Stanford, California 94305
(415) 723-4879

KULIKOWSKI, Casimir, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-2006

LEDERBERG, Joshua, Ph.D.
President
The Rockefeller University
1230 York Avenue
New York, New York 10021
(212) 570-8080, 570-8000

LINDBERG, Donald A.B., M.D. (Past Adv Grp Chrmn)
Director, National Library of Medicine
8600 Rockville Pike
Bethesda, Maryland 20814
(301)496-6221

MYERS, Jack D., M.D.
School of Medicine
Scaife Hall, 1291
University of Pittsburgh
Pittsburgh, Pennsylvania 15261
(412) 648-9933

AIM Advisory Group:

MYERS, Jack D., M.D. (Chairman)
School of Medicine
Scaife Hall, 1291
University of Pittsburgh
Pittsburgh, Pennsylvania 15261
(412) 648-9933

AMAREL, Saul, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-3546

COULTER, Charles L., Ph.D. (Exec. Secretary)
Bldg 31, Room 5B41
Biomedical Research Technology Program
National Institutes of Health
9000 Rockville Pike
Bethesda, Maryland 20892
(301) 496-5411

FEIGENBAUM, Edward A., Ph.D. (Ex-officio)
Co-Principal Investigator - SUMEX
Heuristic Programming Project
Department of Computer Science
701 Welch Road, Building C
Stanford University
Palo Alto, California 94305
(415) 723-4879

KULIKOWSKI, Casimir, Ph.D.
Department of Computer Science
Hill Center Busch Campus
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-2006

LEDERBERG, Joshua, Ph.D.
President
The Rockefeller University
1230 York Avenue
New York, New York 10021
(212) 570-8080, 570-8000

LINDBERG, Donald A.B., M.D.
Director, National Library of Medicine
Building 38, Rm. 2E-17B
8600 Rockville Pike
Bethesda, Maryland 20814
(301) 496-6221

MINSKY, Marvin, Ph.D.
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
545 Technology Square
Cambridge, Massachusetts 02139
(617) 253-5864

MOHLER, William C., M.D.
Associate Director
Division of Computer Research and Technology
National Institutes of Health
Building 12A, Room 3033
9000 Rockville Pike
Bethesda, Maryland 20892
(301) 496-1168

PAUKER, Stephen G., M.D.
Department of Medicine - Cardiology
Tufts New England Medical Center Hospital
171 Harrison Avenue
Boston, Massachusetts 02111
(617) 956-5910

SHORTLIFFE, Edward H., M.D., Ph.D. (Ex-officio)
Principal Investigator - SUMEX
Medical School Office Building, Rm. X271
Stanford University Medical Center
Stanford, California 94305
(415) 723-6979

SIMON, Herbert A., Ph.D.
Department of Psychology
Baker Hall, 339
Carnegie-Mellon University
Schenley Park
Pittsburgh, Pennsylvania 15213
(412) 578-2787, 578-2000

Stanford Community Advisory Committee:

FEIGENBAUM, Edward A., Ph.D. (Chairman)
Heuristic Programming Project
Department of Computer Science
Margaret Jacks Hall
Stanford University
Stanford, California 94305
(415) 723-4879

LEVINTHAL, Elliott C., Ph.D.
Departments of Mechanical and Electrical Engineering
Building 530
Stanford University
Stanford, California 94305
(415) 723-9037

SHORTLIFFE, Edward H., M.D., Ph.D.
Principal Investigator - SUMEX
Medical School Office Building, Rm. X271
Stanford University Medical Center
Stanford, California 94305
(415) 723-6979

Appendix B
Scientific Subproject Abstracts

The following are brief abstracts of our collaborative research projects.

Stanford Project: GUIDON/NEOMYCIN -- KNOWLEDGE ENGINEERING FOR TEACHING MEDICAL DIAGNOSIS

Principal Investigators: William J. Clancey, Ph.D.
701 Welch Road
Department of Computer Science
Stanford University
Palo Alto, California 94304
(415) 723-1997 (CLANCEY@SUMEX-AIM)

Bruce G. Buchanan, Ph.D.
Computer Science Department
701 Welch Road
Stanford University
Palo Alto, California 94304
(415) 723-0935 (BUCHANAN@SUMEX-AIM)

SOFTWARE AVAILABLE ON SUMEX

GUIDON--A system developed for intelligent computer-aided instruction. Although it was developed in the context of MYCIN's infectious disease knowledge base, the tutorial rules will operate upon any EMYCIN knowledge base.

NEOMYCIN--A consultation system derived from MYCIN, with the knowledge base greatly extended and reconfigured for use in teaching. In contrast with MYCIN, diagnostic procedures, common sense facts, and disease hierarchies are factored out of the basic finding/disease associations. The diagnostic procedures are abstract (not specific to any problem domain) and model human reasoning, unlike the exhaustive, top-down approach implicit in MYCIN's medical rules. This knowledge base will be used in the GUIDON2 family of instructional programs, being developed on D-machines.

REFERENCES

Clancey, W.J.: **Knowledge-Based Tutoring: The GUIDON Program**, Cambridge: The MIT Press, 1987.

Clancey, W.J.: *Methodology for building an intelligent tutoring system.* In Kintsch, Polson, and Miller, (Eds.), METHODS AND TACTICS IN COGNITIVE SCIENCE. L. Erlbaum Assoc., Hillsdale, NJ. 1984. (Also STAN-CS-81-894, HPP 81-18)

Clancey, W.J.: *Acquiring, representing, and evaluating a competence model of diagnosis.* In Chi, Glaser, and Farr (Eds.), THE NATURE OF EXPERTISE. In preparation. HPP-84-2.

Stanford Project: MOLGEN -- AN EXPERIMENT PLANNING SYSTEM FOR MOLECULAR GENETICS

Principal Investigators: Edward A. Feigenbaum, Ph.D.
Department of Computer Science
Stanford University

Charles Yanofsky, Ph.D. (YANOFSKY@SUMEX-AIM)
Department of Biology
Stanford University
Stanford, California 94305
(415) 725-3815

Contact: Dr. Peter FRIEDLAND@SUMEX-AIM
(415) 723-3728

The MOLGEN project has focused on research into the applications of symbolic computation and inference to the field of molecular biology. This has taken the specific form of systems which provide assistance to the experimental scientist in various tasks, the most important of which have been the design of complex experiment plans and the analysis of nucleic acid sequences. Our current research concentrates on scientific discovery within the subdomain of regulatory genetics. We desire to explore the methodologies scientists use to modify, extend, and test theories of genetic regulation, and then emulate that process within a computational system.

Theory or model formation is a fundamental part of scientific research. Scientists both use and form such models dynamically. They are used to predict results (and therefore to suggest experiments to test the model) and also to explain experimental results. Models are extended and revised both as a result of logical conclusions from existing premises and as a result of new experimental evidence.

Theory formation is a difficult cognitive task, and one in which there is substantial scope for intelligent computational assistance. Our research is toward building a system which can form theories to explain experimental evidence, can interact with a scientist to help to suggest experiments to discriminate among competing hypotheses, and can then revise and extend the growing model based upon the results of the experiments.

The MOLGEN project has continuing computer science goals of exploring issues of knowledge representation, problem-solving, discovery, and planning within a real and complex domain. The project operates in a framework of collaboration between the Heuristic Programming Project (HPP) in the Computer Science Department and various domain experts in the departments of Biochemistry, Medicine, and Biology. It draws from the experience of several other projects in the HPP which deal with applications of artificial intelligence to medicine, organic chemistry, and engineering.

SOFTWARE AVAILABLE ON SUMEX

SPEX system for experiment design.
UNITS system for knowledge representation and acquisition.
SEQ system for nucleotide sequence analysis.

REFERENCES

1. Friedland, P.E.: *Knowledge-based experiment design in molecular genetics*, (Ph.D. thesis). Stanford Computer Science Report, STAN-CS-79-771.

2. Friedland, P.E. and Iwasaki, Y.: *The concept and implementation of skeletal plans,* Journal of Automated Reasoning, 1(2):161-208, 1985.

3. Friedland, P.E. and Kedes, L.: *Discovering the secrets of DNA,* Communications of the ACM, 28(11):1164-1186, November, 1985.

4. Stefik, M.J.: *An examination of a frame-structured representation system,* Proc. Sixth IJCAI, Tokyo, August, 1979, pp. 845-852.

5. Stefik, M.J.: *Planning with constraints,* (Ph.D. thesis). Stanford Computer Science Report, STAN-CS-80-784, March, 1980.

6. Karp, P., and D. Wilkins: *An Analysis of the Deep/Shallow Distinction for Expert Systems.* Stanford University Knowledge Systems Laboratory Report KSL-86-32, 1986.

7. Karp, P., and P. Friedland: *Coordinating the Use of Qualitative and Quantitative Knowledge in Declarative Device Modeling.* Stanford University Knowledge Systems Laboratory Report KSL-87-09, 1987.

8. Round, A.: *QSOPS: A Workbench Environment for the Qualitative Simulation of Physical Processes.* Stanford University Knowledge Systems Laboratory Report KSL-87-37, 1987.

Stanford Project: ONCOCIN -- KNOWLEDGE ENGINEERING FOR ONCOLOGY CHEMOTHERAPY CONSULTATION

Principal Investigator: Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University Medical Center
Medical School Office Building
Stanford, California 94305
(415) 723-6979 (SHORTLIFFE@SUMEX-AIM)

Project Director: Dr. Lawrence M. Fagan (FAGAN@SUMEX-AIM)

The ONCOCIN Project is overseen by a collaborative group of physicians and computer scientists who are developing an intelligent system that uses the techniques of knowledge engineering to advise oncologists in the management of patients receiving cancer chemotherapy. The general research foci of the group members include knowledge acquisition, inexact reasoning, explanation, and the representation of time and of expert thinking patterns. Much of the work developed from research in the 1970's on the MYCIN and EMYCIN programs, early efforts that helped define the group's research directions for the coming decade. MYCIN and EMYCIN are still available on SUMEX for demonstration purposes.

The prototype ONCOCIN system is in limited experimental use by oncologists in the Stanford Oncology Clinic. Thus, much of the emphasis of this research has been on human engineering so that the physicians will accept the program as a useful adjunct to their patient care activities. ONCOCIN has generally been well-accepted since its introduction, and we are now testing a version of the program which runs on professional workstations (rather than the central SUMEX computer) so that it can be implemented and evaluated at sites away from the University.

SOFTWARE AVAILABLE ON SUMEX

MYCIN-- A consultation system designed to assist physicians with the selection of antimicrobial therapy for severe infections. It has achieved expert level performance in formal evaluations of its ability to select therapy for bacteremia and meningitis. Although MYCIN is no longer the subject of an active research program, the system continues to be available on SUMEX for demonstration purposes and as a testing environment for other research projects.

EMYCIN-- The "essential MYCIN" system is a generalization of the MYCIN knowledge representation and control structure. It is designed to facilitate the development of new expert consultation systems for both clinical and non-medical domains.

ONCOCIN-- This system is in clinical use but requires Lisp machines to be run. Much of the knowledge in the domain of cancer chemotherapy is already well-specified in protocol documents, but expert judgments also need to be understood and modeled.